Argonne NATIONAL LABORATORY

Introductory Course: Using LS-OPT® on the TRACC Cluster

2.1 – Best Practices on the TRACC Cluster

By: Dr.-Ing. Hubert Ley

Transportation Simulation Project Leader

U.S. DEPARTMENT OF ENERGY

TRACC - High Performance Computing for Transportation Research and Applied Technology

- USDOT and USDOE transportation research programs, private industry and state and regional transportation agencies are moving to simulation based design and analysis for improvements in efficiency, economics and safety
- Higher fidelity analysis in areas such as crashworthiness, aerodynamics, combustion, thermal management, weather modeling and traffic simulation require access to state-of-the-art computational and visualization facilities
- ANL has the necessary expertise in high performance computing and transportation system analysis to provide both a national HPC user facility and a focal point for computational research for transportation applications

About Argonne: One of DOE's Largest Research Facilities

http://www.anl.gov/

- Located 25 miles from the Chicago Loop, it was the first national laboratory, chartered in 1946
- Operated by the University of Chicago for the U.S. Department of Energy
- Major research missions include basic science, environmental management, and advanced energy technologies
- About 3,000 employees, including about 1,000 scientists and engineers, of whom 750 hold doctorate degrees
- Annual operating budget of about $475 million (80% from DOE)
- Since 1990, Argonne has worked with more than 600 companies and numerous federal agencies.

TRACC Is Being Built as a National DOT Supercomputer Facility

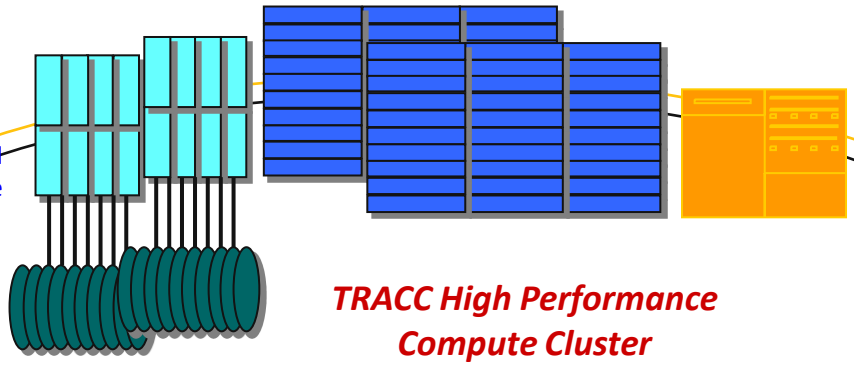

180TB Global Parallel File System Disk Storage

LTO-3 160TB Archive/Backup Tape Storage

TRACC High Performance Compute Cluster

1024core /1 28 compute nodes

(two quad-core AMD Opteron CPUs 8GB of RAM)

High-bandwidth connectivity to the Illinois Wired/Wireless Infrastructure for Research and Education (I-WIRE) and Internet2.

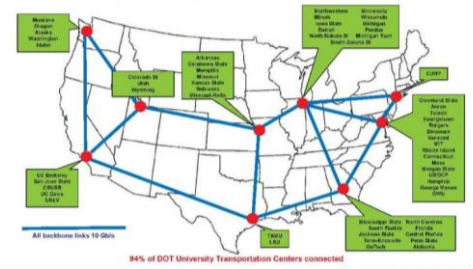

TRACC Training Center

Argonne Nuclear Engineering Division Linux Cluster 10Gb Dedicated Access

Argonne LCRC JAZZ Linux Cluster 10Gb Dedicated Access

Internet2 Network

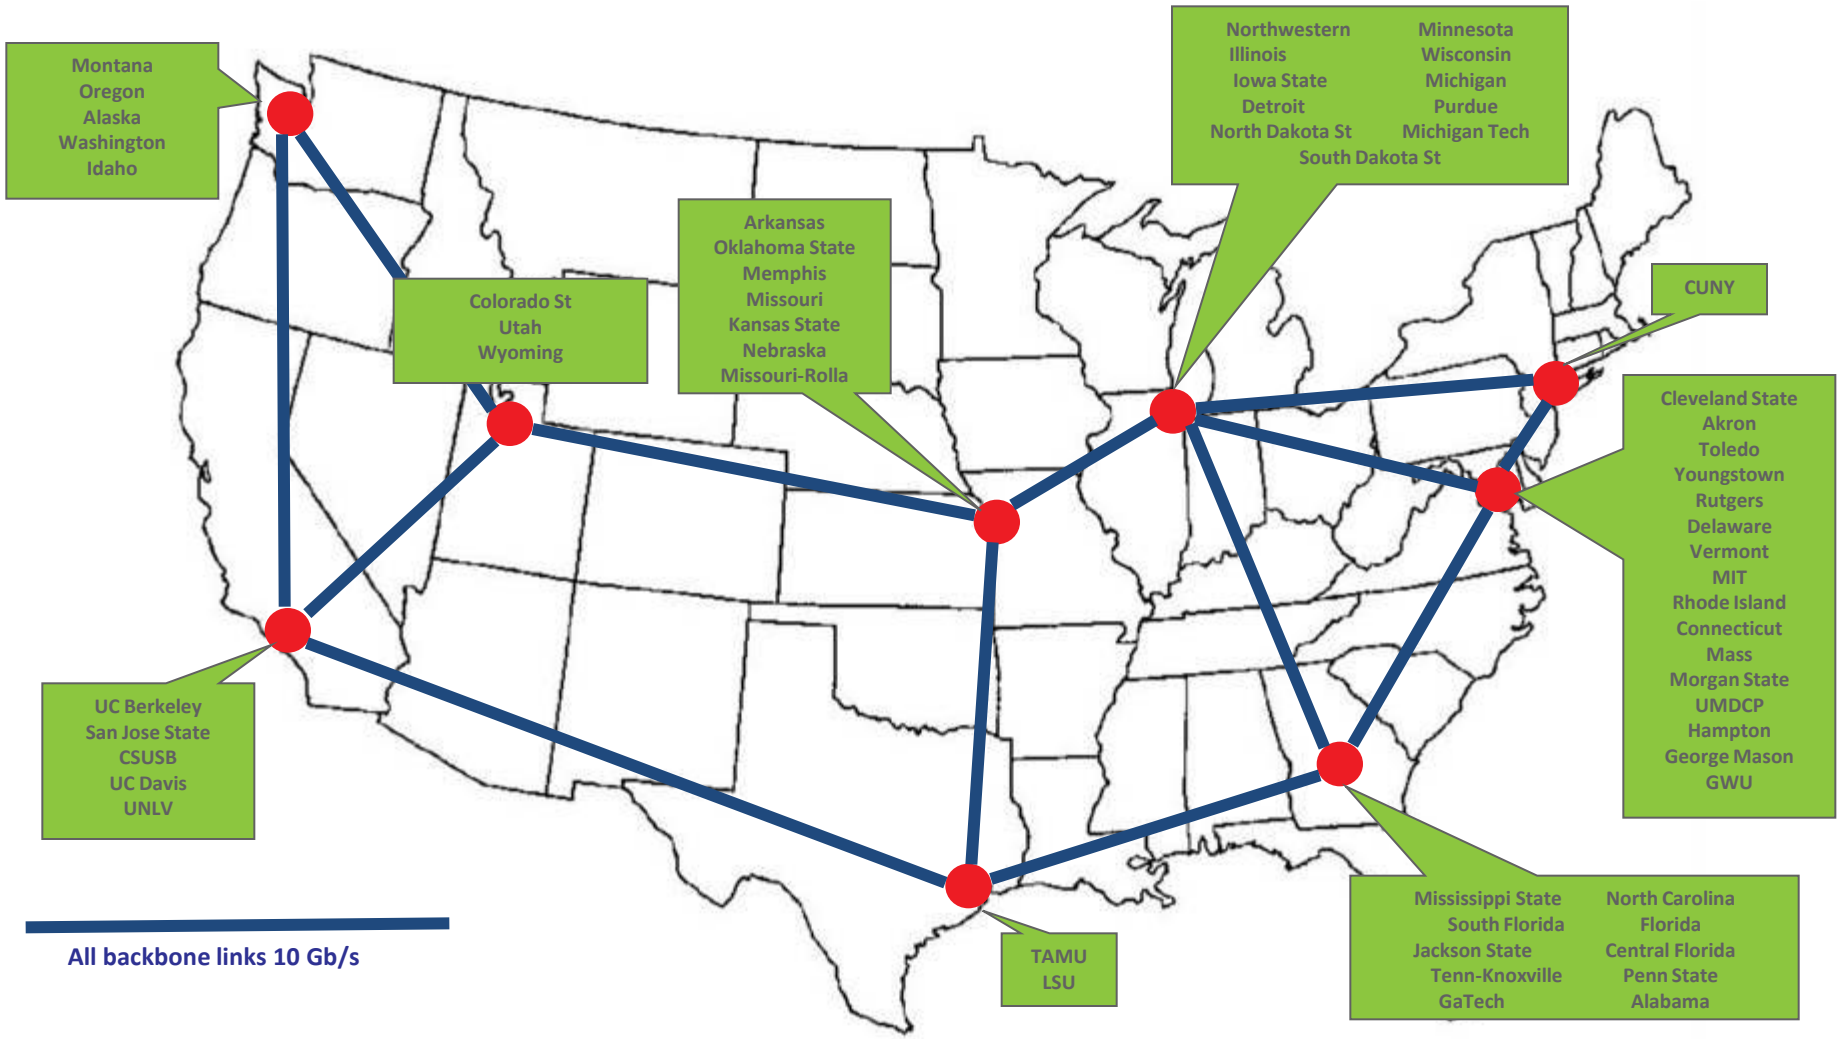

94% of DOT University Transportation Centers connected

TRACC Is a National USDOT Supercomputing Facility

TRACC High Performance Compute Cluster

- *1024 core /128 compute nodes*
- 180TB Global Parallel File System Disk Storage
- 160TB Archive/Backup Tape Storage

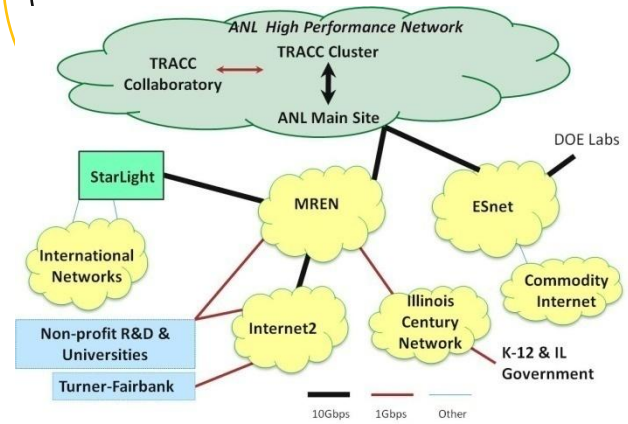

High-bandwidth connectivity is provided via the Argonne high-performance network to world-wide research and education networks (Internet2 and ESnet)

TRACC Collaboratory - Visualization, Access Grid, and Digital Conferencing

TRACC HPC Configuration Diagram

The TRACC computational cluster is a customized LS-1 system from Linux Networx consisting of 1024 core 128 compute nodes, each with two quad-core AMD Opteron CPUs and 8GB of RAM, a DataDirect Networks storage system consisting of 180TB of shared RAID storage, expandable to 750TB, a high-bandwidth, low-latency InfiniBand network for computations, and a high-bandwidth Gigabit Ethernet management network. The system will also include the highest-performance compiler and MPI library available for the AMD Opteron architecture. with a peak performance of ~4 TFlops

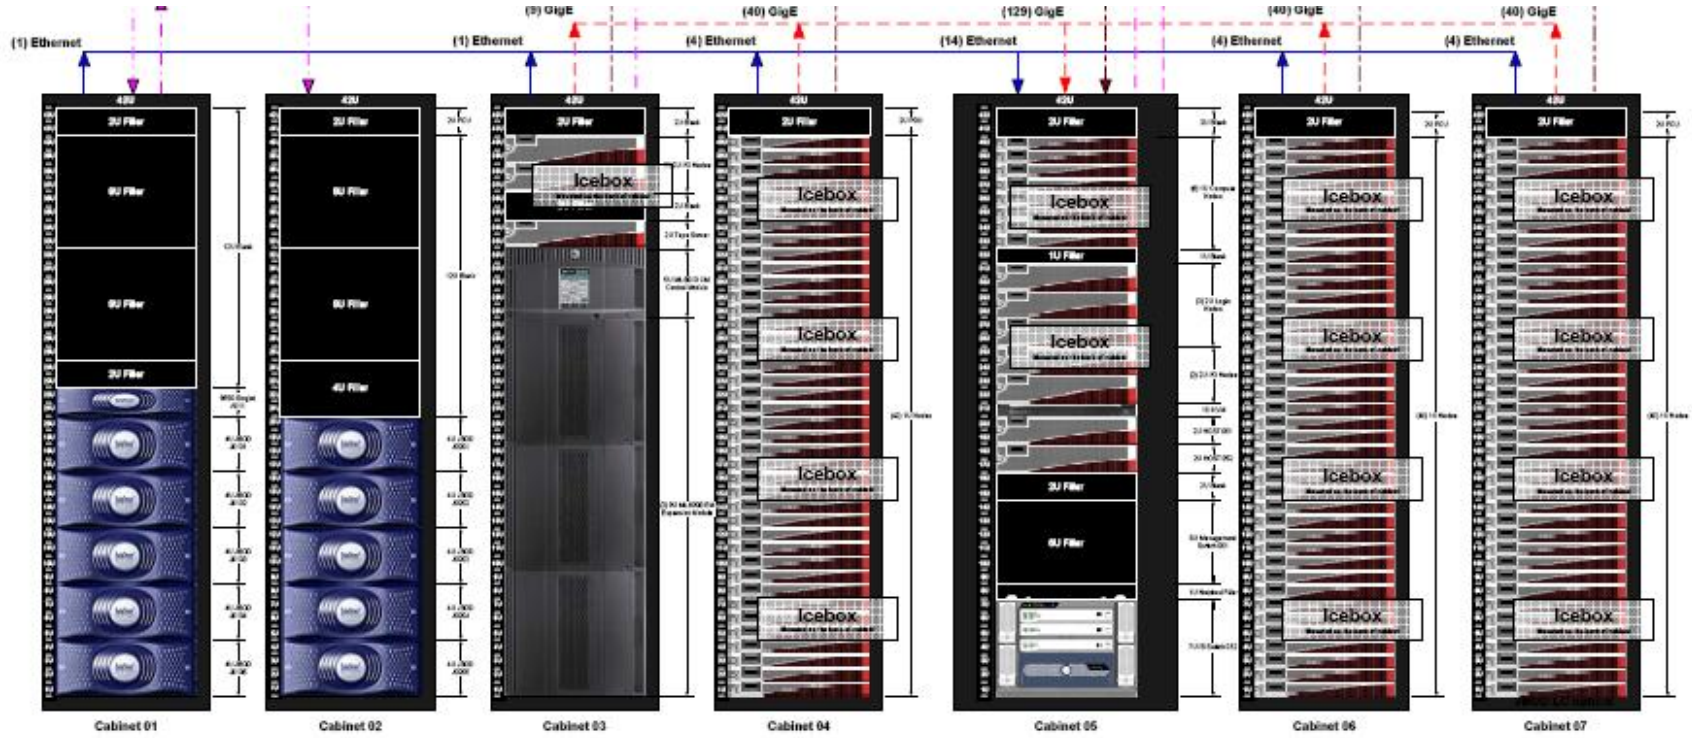

Other Applications: Computational Fluid Dynamics

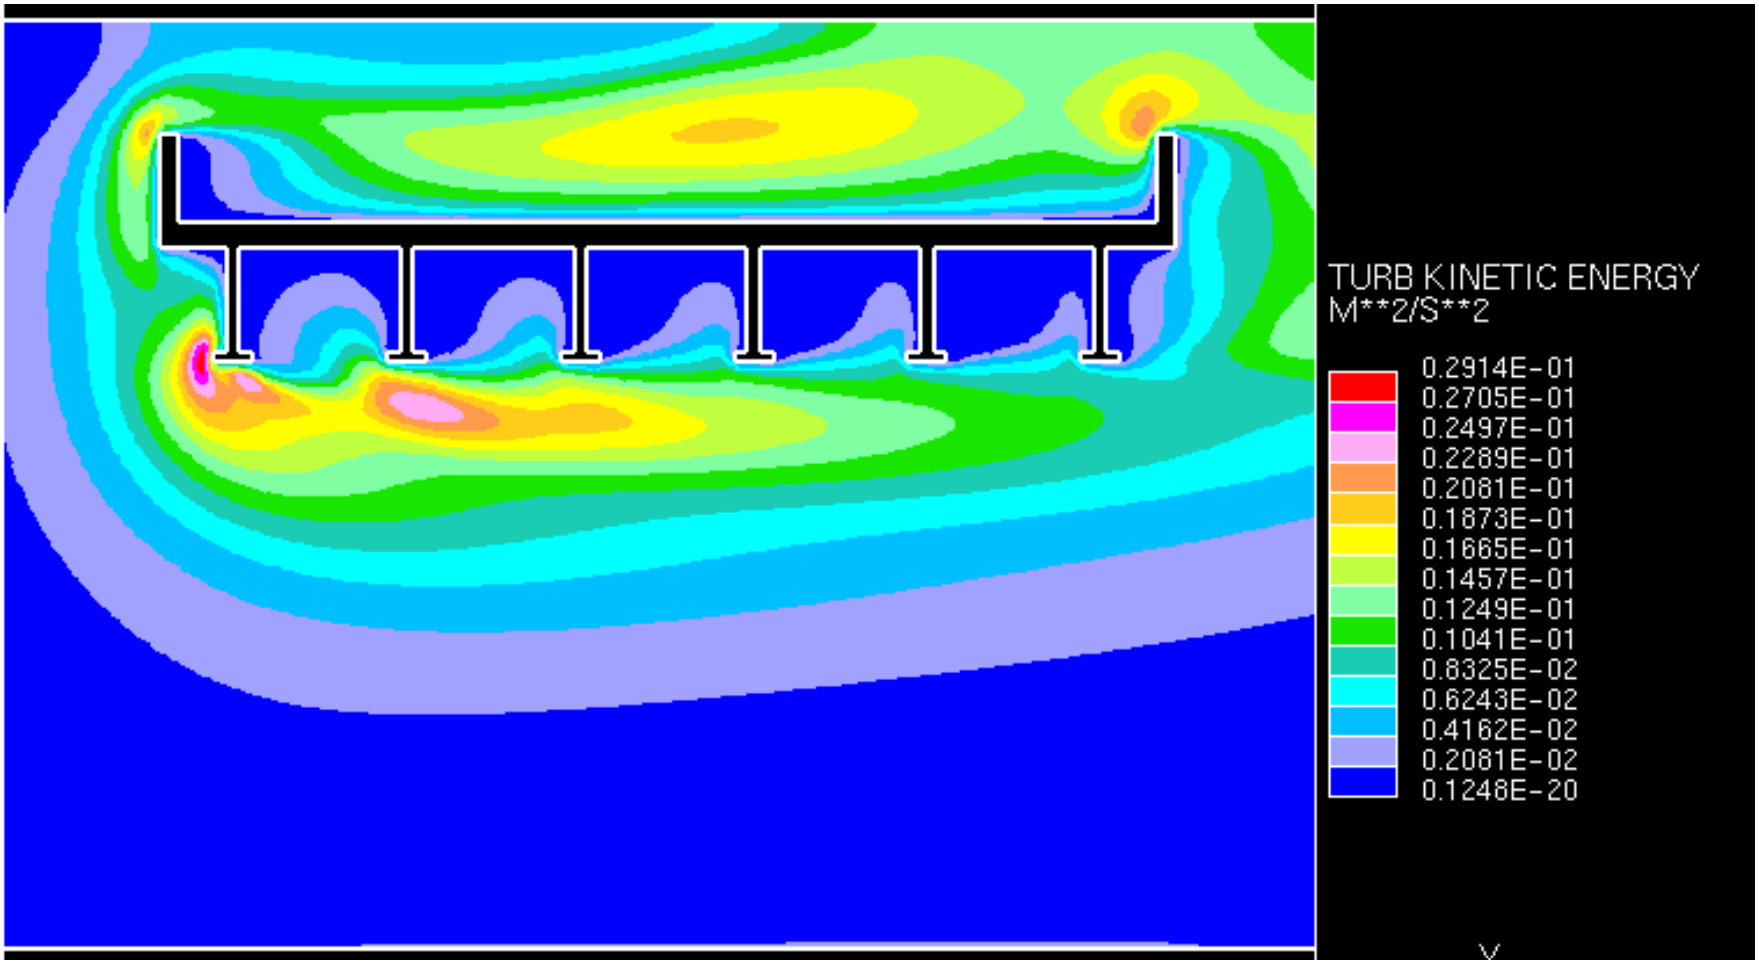

Other Applications: Transportation Simulation

- TRANSIMS has unique capabilities to simulate the effects of emergency evacuations.
- This project has been implemented to model the effects of a no-notice event on the multi-modal regional transportation system in the Chicago metropolitan area.
- This project deals with the dynamic effect on the transportation system.
- Sponsored by the Illinois Terrorism Task Force, the Illinois Department of Transportation, the City of Chicago, and other federal, state and local agencies.

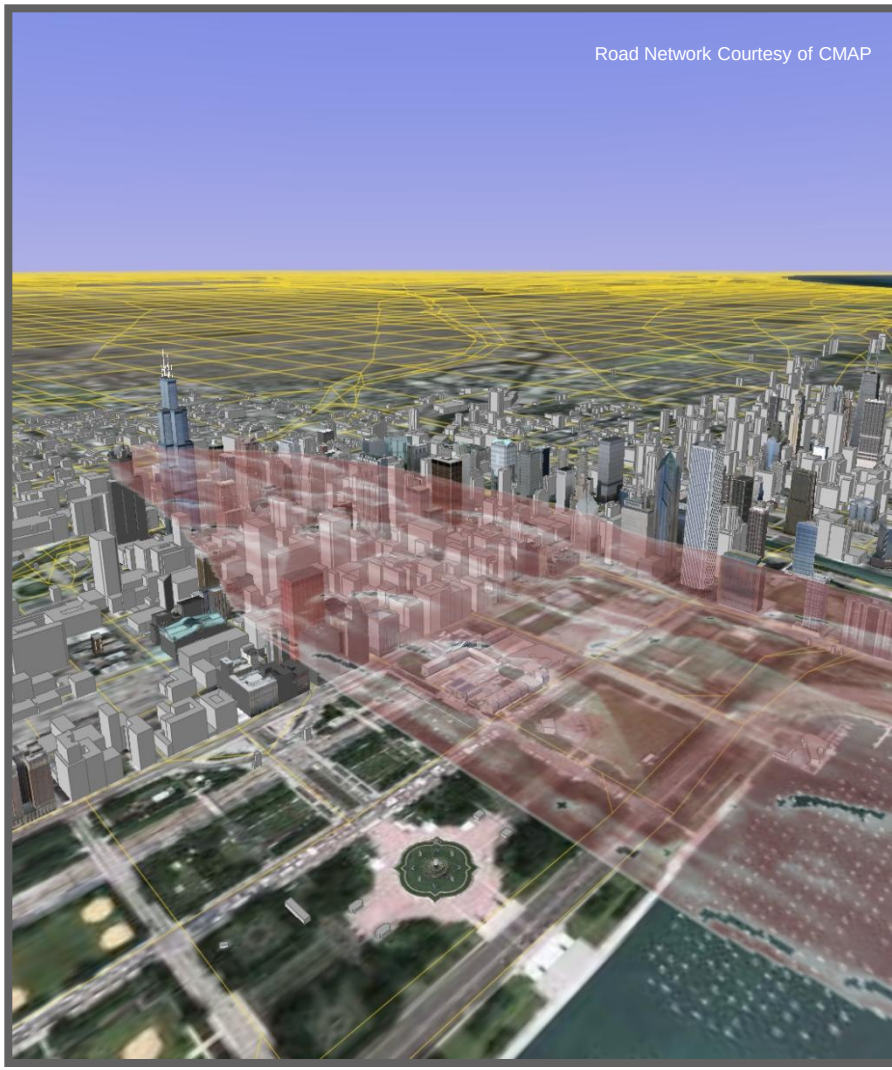

Security Considerations

- SSH is at this time the only protocol provided for accessing the cluster remotely
- Accounts and passwords can be obtained through:
  - http://www.anl.gov/TRACC/Users/becoming.html
- Accounts are assigned under the rules of the Department of Energy
- Accounts are not allowed to be shared with anyone – sharing an account and password will result in a permanent loss of user privileges
- Don't expect any privacy when using the machine – this is a federal computing system and therefore subject to many restrictions

Signup and Allocation

- The proposed project must be a large-scale, complex application of either national or regional transportation significance
- Massively parallel computing and visualization resources for the project are not readily available elsewhere
- The project should provide education, technology transfer, or training (capacity building for national, regional, state, and local transportation practioners)

Transportation Research and Analysis Computing Center
Request for Allocation of Services

Please complete the below information and send it to TRACC at: TRACC-Help@anl.gov, or fax it to: Karen Ley – 630-578-4257.

Project Name:		
Focus Area		
TRACC Contact/Sponsor		
Requester Contact Information:	Name:	
	Title:	
	Organization:	
	Address:	
	Telephone:	
	Fax:	
	Email:	
Cognizant Engineer or Sponsor: (USDOT, State DOT, TRB, ASSCHTO, USDOT Cooperative Research Programs:	Name:	
	Title:	
	Organization:	
	Address:	
	Telephone:	
	Fax:	
	Email:	

Conditions of Use

Security Notice

- This web site is part of a Federal computer system used to accomplish Federal functions. The Department of Energy monitors this web site for security purposes to ensure it remains available to all users and to protect information in the system. By accessing this web site, you are expressly consenting to these monitoring activities. Unauthorized attempts to defeat or circumvent security features, to use the system for other than intended purposes, to deny service to authorized users, to access, obtain, alter, damage, or destroy information, or otherwise to interfere with the system or its operation are prohibited. Evidence of such acts may be disclosed to law enforcement authorities and result in criminal prosecution under the Computer Fraud and Abuse Act of 1986 (Pub. L. 99-474) and the National Information Infrastructure Protection Act of 1996 (Pub. L. 104-294), (18 U.S.C. 1030), or other applicable criminal laws.

The TRACC Cluster - Design

- All individual machines are based on the same hardware to increase reliability and compatibility across the system:
  - CPUs: Two Quad-Core AMD Opteron(tm) Processors
  - Clocking Speed: 2.4GHz, Memory: 8GB of RAM
- 3 Login Servers provide user-level access to the cluster
  - Round robin assignment on a single public address, 8GB of memory
- 128 Computing Nodes with 8GB available through PBS job scheduling
- 2 Administrative Servers for cluster management
- 4 File Servers to provide a shared GPFS file system
- 1 Tape Server for backup services
- A 180GB shared RAID file system with 480 x 500G SATA drives
- A 160GB tape capacity robot
- InfiniBand interconnectivity between servers and computing nodes
- Additional high speed Ethernet for management
- Performance is approximately 4 Tera Flops

Access to the TRACC Cluster

- Access to the TRACC cluster is provided primarily by SSH (secure shell) protocol
- Suggested Client Software:
  - SSH Secure Shell (www.ssh.com) for Windows (terminal access)
  - PuTTY (www.putty.org) for Windows (terminal access)
  - Standard ssh client software under Linux
  - WinSCP (winscp.net) for Windows (for efficient file access)
- Login is provided under a single address (round robin assignment to one of the three login servers):

login.tracc.anl.gov

- To use the cluster, an ANL domain account or collaborator account is required

http://www.anl.gov/TRACC/Users/becoming.html

Terminal Sessions

- SSH Secure Shell
  - Allows for terminal (tty) shell access
  - Provides a file manager for uploads and downloads
  - Commercial version is available

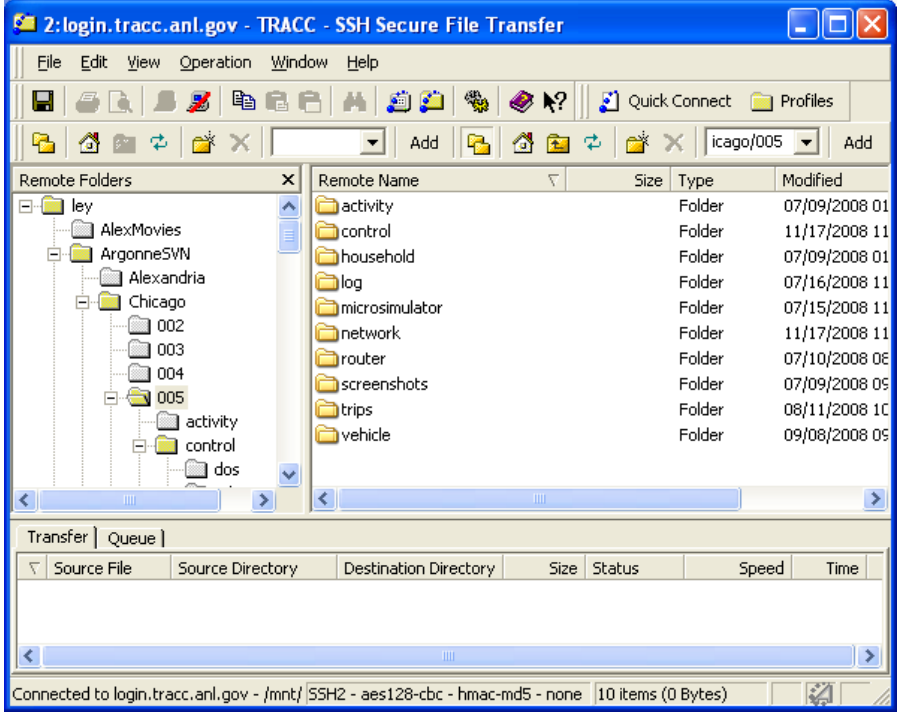

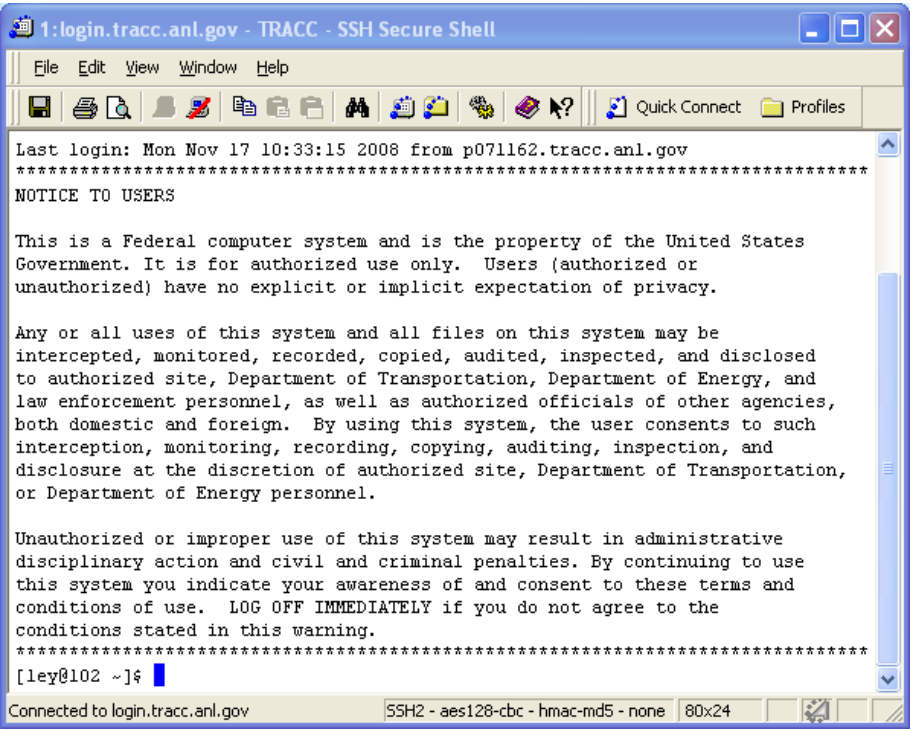

- Commercial Version at
  - www.ssh.com
- Comments on the non-commercial version at:
  - http://www.ssh.com/support/downloads/secureshellwks/non-commercial.html

Terminal Sessions

- PuTTY
  - Open source software
  - Allows for terminal (tty) shell access
  - Authentication manager is available to cache private keys for multiple concurrent sessions
  - Does not support file downloading and uploading directly

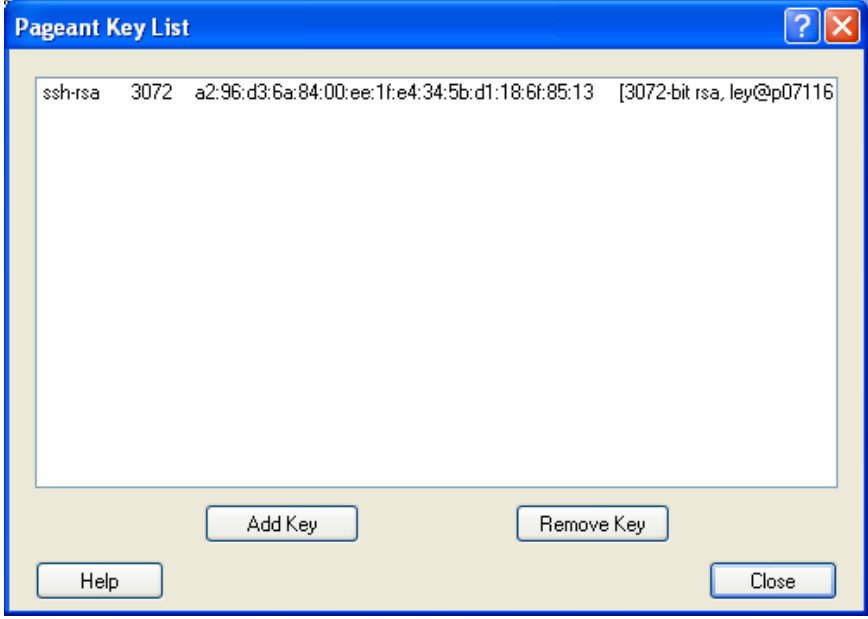

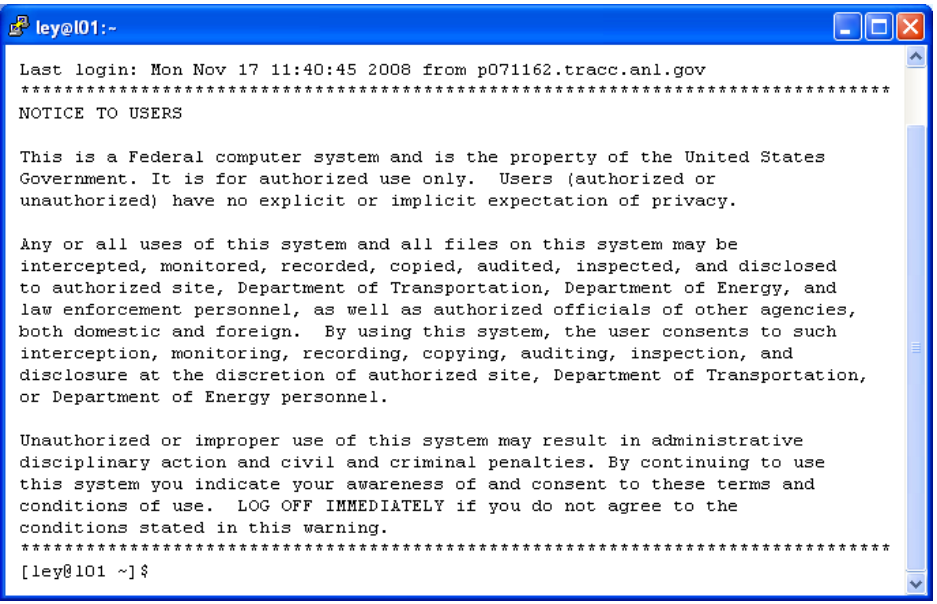

- Downloadable from
  - http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html
  - PuTTY can be run as a user application without installation
  - This is non-commercial software and can be used freely (under the conditions outlined in the license)

Terminal Sessions and Remote File Managers

- WinSCP
  - File manager functionality only, no terminal sessions
  - But, WinSCP interoperates easily with PuTTY
  - Allows better integration with Windows applications than SSH tools
- Free software from
  - winscp.net
- Sourceforge is an open source project

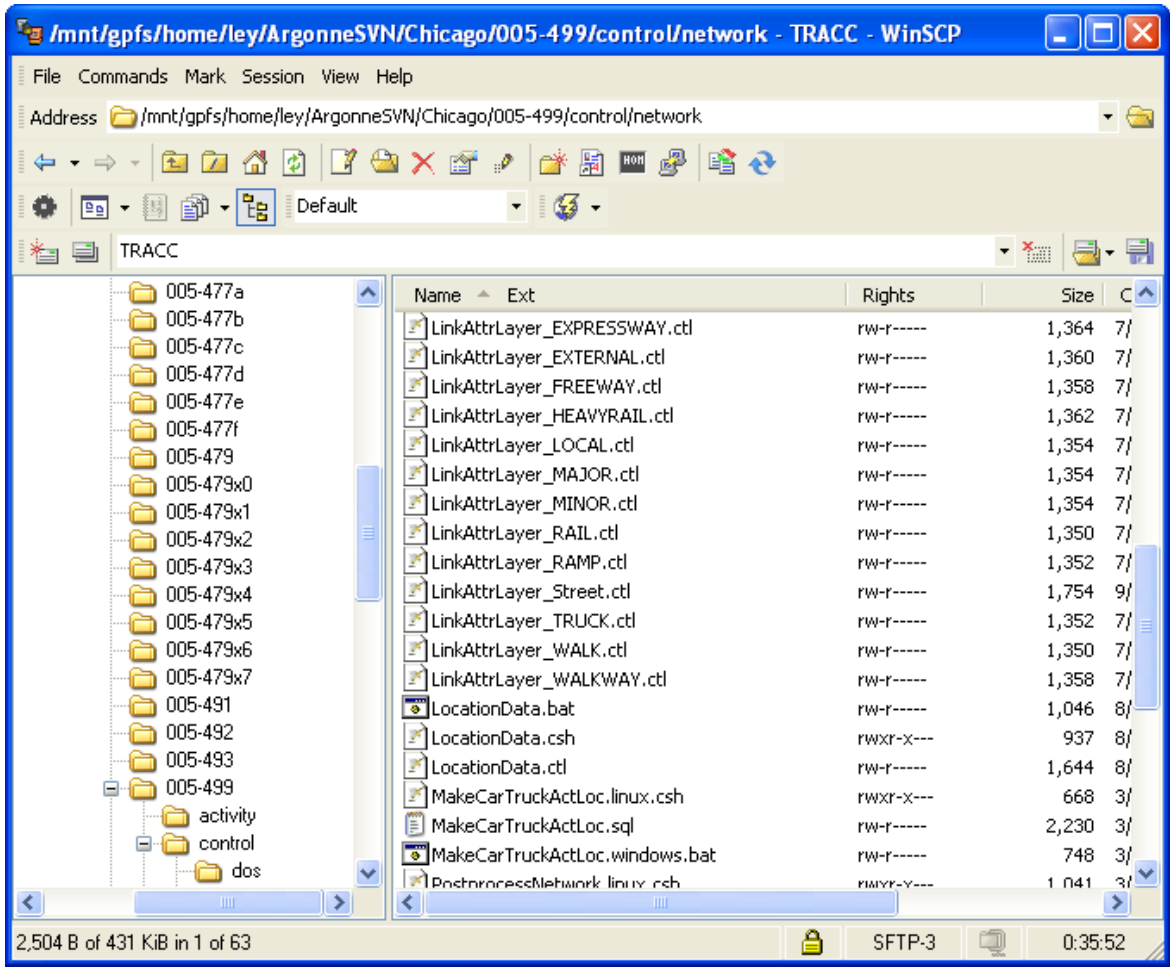

Graphical Terminal Sessions

- NoMachine NX
  - Full screen Linux desktop
  - Efficient, secure, and fast
- www.nomachine.com
  - NX Client can be used to establish a complete desktop connection

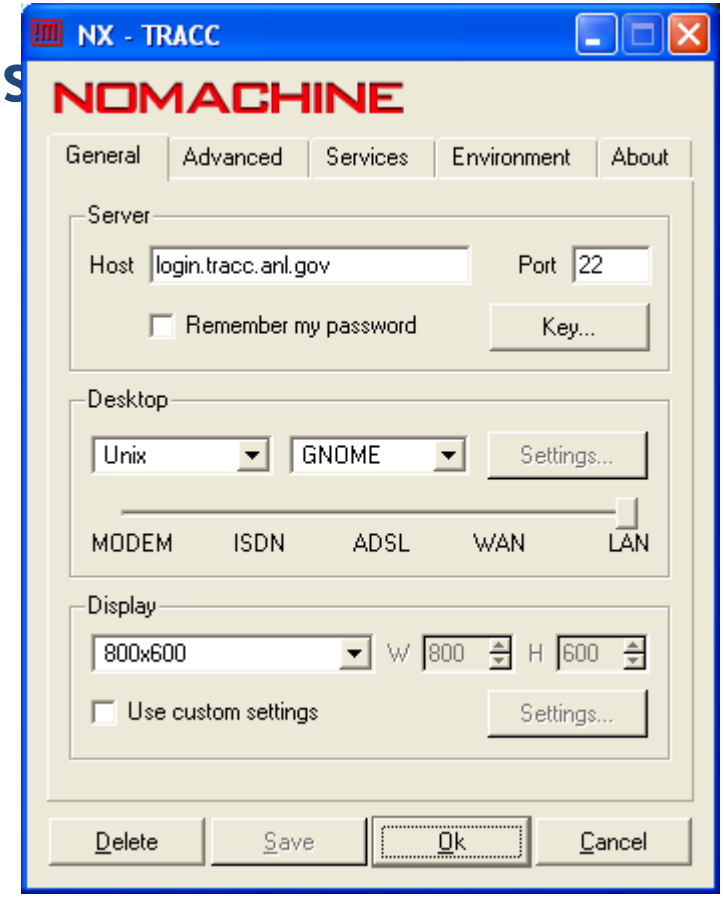

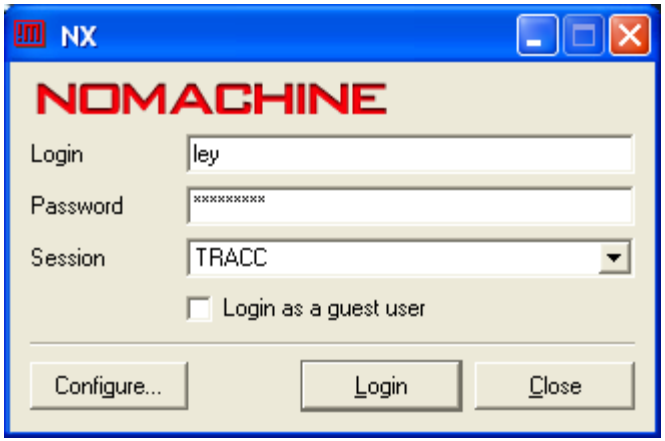

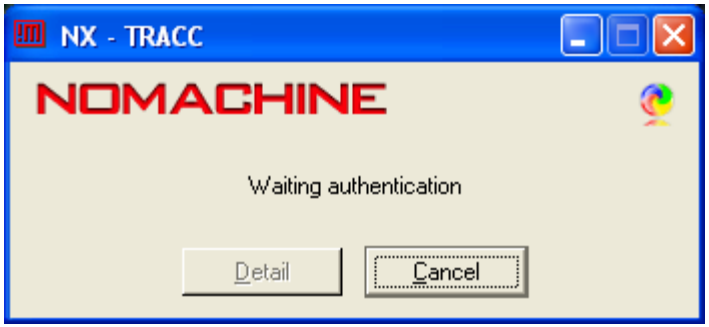

Graphical Terminal Sessions

- Access to a complete Linux desktop
- Compressed data transfer
- Optimized for slow data connections, e.g. DSL
- Based on X protocol
- Can be used for nearly any Linux application

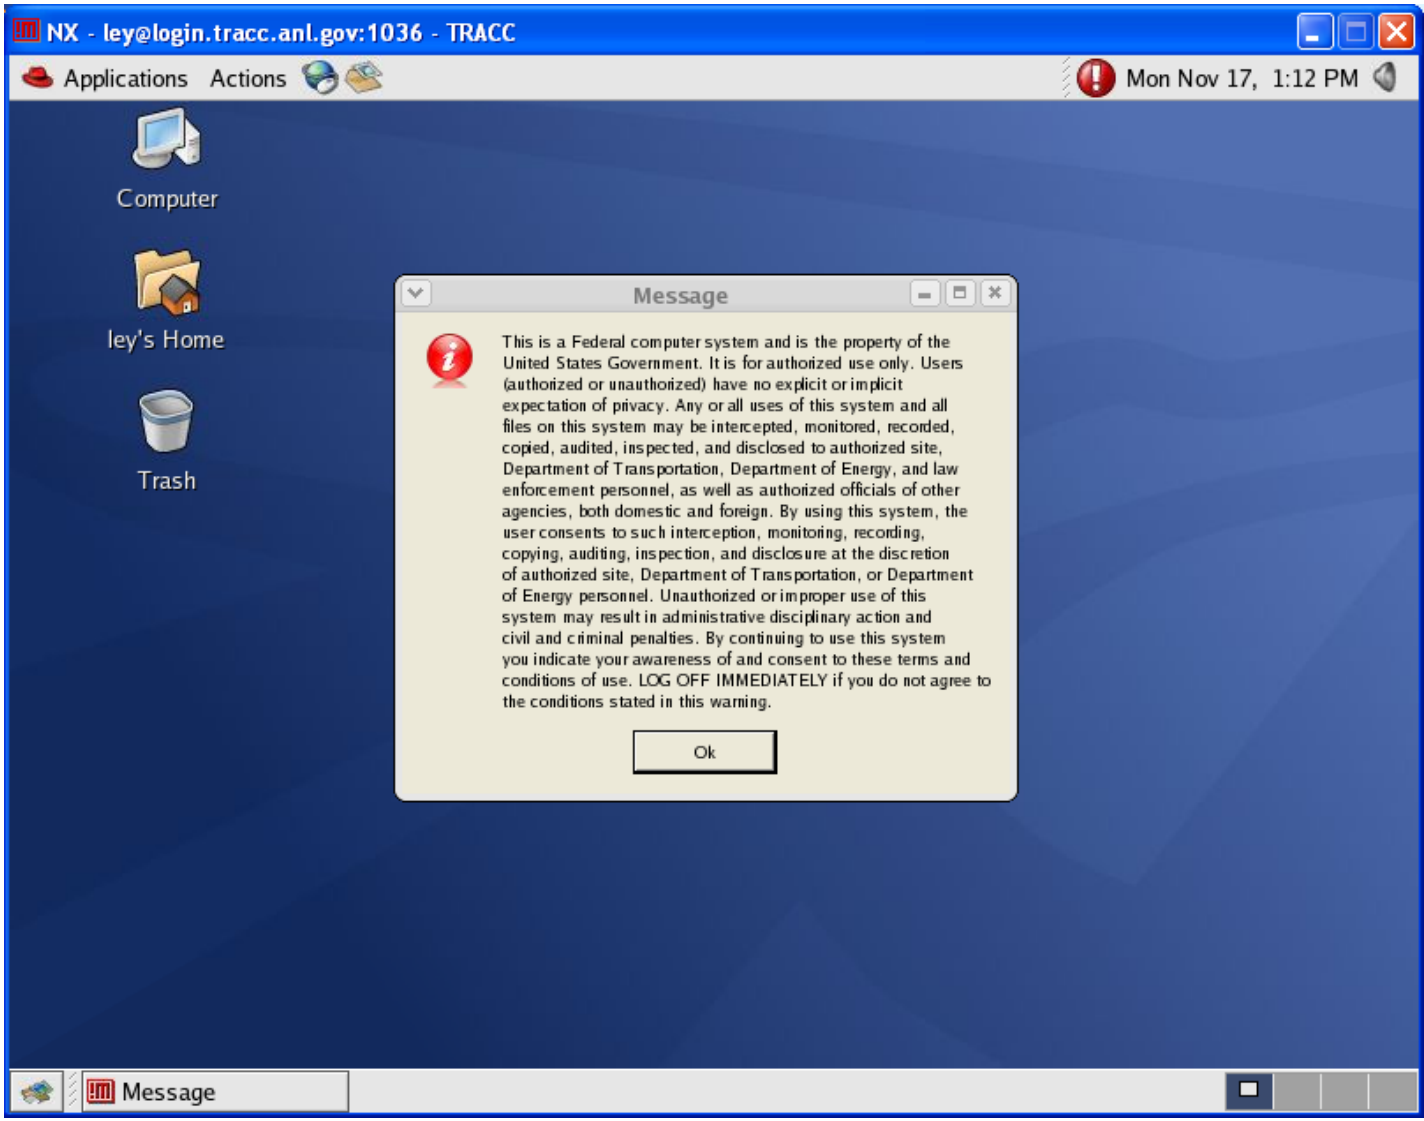

Job Control and Job Submission

- The three login servers are primarily used to access the cluster, to edit files, to prepare data, and to compile codes
  - The login nodes are shared by all users concurrently using the cluster
  - Each login node may provide access to dozens of users concurrently
  - No jobs should be run on these machines that require significant CPU time or memory

- The cluster nodes are assigned to run jobs through a job scheduling system (PBS – Portable Batch System)
  - Submission of jobs using “qsub”
  - Status inquiries using “qstat”
  - Details using “tracejob”

- Instructions can be found at:
  - https://wiki.anl.gov/tracc/Using_the_cluster

TRACC User Support and Wiki

- Access to the cluster is currently restricted to pre-approved hosts. Please contact us with the static hostnames or IP addresses of the systems from which you would like to access the cluster.

- Under Windows, you can find out your host name and IP address with the command "ipconfig". Note that IP addresses beginning with "10." or "192.168." are private addresses, and are not accessible outside your own network. If you have such an address, you can find out your public address at (for example) http://whatismyip.com

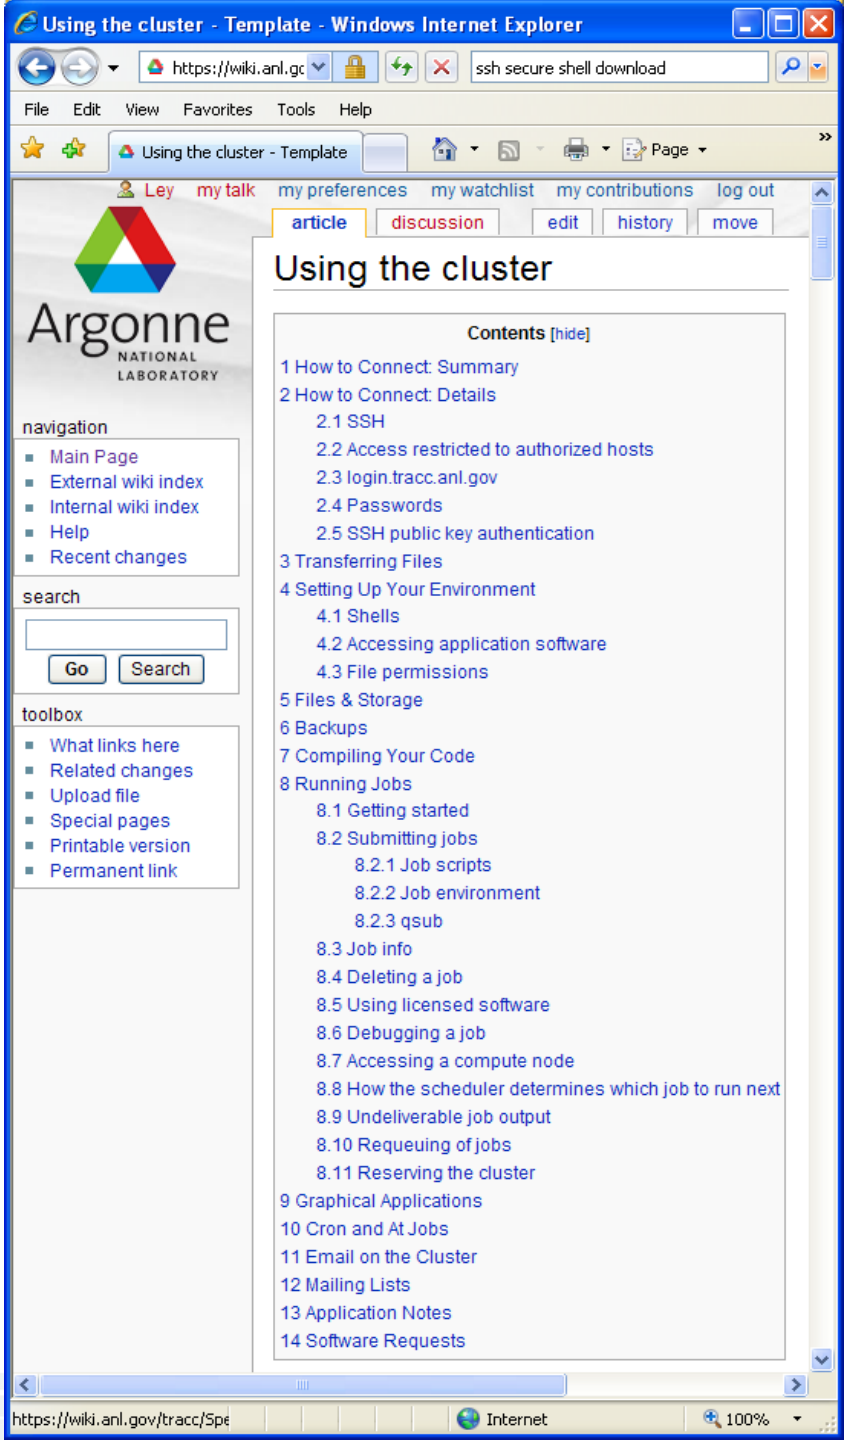

Credits and Acknowledgements

- USDOT provided the funding for the development of these training materials
- USDOT provided the funding for the TRACC computing center and the resources necessary to perform these training session